

CLAIMS

What is claimed is:

1. A method for providing convergence of data copies in asynchronous data replication in a database system, the database system including a plurality of nodes with a plurality of table copies, comprising:

- (a) labeling rows of the plurality of table copies with a monotonic number, a copy identification, and propagation controls;
- (b) asynchronously capturing at least one labeled change to any row of any of the plurality of table copies from a database recovery log;
- (c) determining that the captured change is to be communicated to others of the plurality of table copies;
- (d) communicating the captured change to the others of the plurality of table copies; and
- (e) applying the communicated change to the others of the plurality of table copies, wherein the plurality of table copies converge to a same state.

15 2. The method of claim 1, wherein the monotonic number comprises a non-decreasing time value, wherein the time values for each of the plurality of table copies are in a common time base.

20 3. The method of claim 2, wherein the monotonic number is automatically applied to a changed table copy in response to a user induced change.

4. The method of claim 1, wherein the copy identification is assigned to each of the plurality of table copies, wherein the copy identification uniquely identifies each of the plurality of table copies, wherein the copy identifications for each of the plurality of table copies have an ordering property.

5

5. The method of claim 4, wherein the copy identification is automatically applied to a changed table copy in response to a user induced change.

10 6. The method of claim 1, wherein the propagation controls comprise a delete

label and a conflict label.

15 7. The method of claim 6, wherein the delete label indicates that a row delete in any of the plurality of table copies is not to be communicated to the others of the plurality of table copies.

15

8. The method of claim 6, wherein the conflict label indicates that an implicit row delete in any of the plurality of table copies is to be communicated to the others of the plurality of table copies.

20

9. The method of claim 1, wherein the capturing (b) comprises:

- (b1) accessing at least one entry of the database recovery log pertaining to the labeled change;
- (b2) determining a type of the labeled change;

- (b3) extracting old column values and new column values of the labeled change from the database recovery log entry; and
- (b4) extracting changed key column values and unchanged key column values from the database recovery log entry.

5

10. The method of claim 9, wherein the type of the labeled change comprises a row insert, a row delete, a non-key update, or a key update.

11. The method of claim 1, wherein the determining (c) comprises:

10 (c1) examining a labeled change type, the copy identification, and the propagation control values for the captured change, wherein the propagation control values comprises a delete label and a conflict label.

15 12. The method of claim 11, wherein the determining (c) further comprises:

(c2) determining that the captured change is to be communicated to the others of the plurality of table copies if the labeled change type is a row insert and if a copy identification for the captured change is a copy identification assigned to a current table copy.

20 13. The method of claim 11, wherein the determining (c) further comprises:

(c2) determining that the captured change is to be communicated to the others of the plurality of table copies if the labeled change type is a row delete and if the delete label of the captured change indicates that the captured change is to be communicated.

14. The method of claim 11, wherein the determining (c) further comprises:

(c2) determining that the captured change is to be communicated to the others of
the plurality of table copies if the labeled change type is a non-key update or a key update
and if neither the delete label nor the conflict label of the captured change indicates that the
captured change is not to be communicated, and if a new copy identification of the captured
change is a copy identification assigned to a current table copy.

5

15. The method of claim 1, wherein the communicating (d) comprises:

(d1) sending information for the captured change to the others of the plurality of
table copies, wherein if the captured change is a row insert, the information comprises key
column values, non-key column values, a monotonic number, and a copy identification of
the captured change.

10

15

16. The method of claim 1, wherein the communicating (d) comprises:

(d1) sending information for the captured change to the others of the plurality of
table copies, wherein if the captured change is a row delete, the information comprises key
column values, a monotonic number, and a copy identification of the captured change.

20

17. The method of claim 1, wherein the communicating (d) comprises:

(d1) sending information for the captured change to the others of the plurality of
table copies, wherein if the captured change is a non-key update, the information comprises
key column values, new non-key column values, an old monotonic number, a new
monotonic number, an old copy identification, and a new copy identification of the captured

change.

18. The method of claim 1, wherein the communicating (d) comprises:

(d1) sending information for the captured change to the others of the plurality of

5 table copies, wherein if the captured change is a key update, the information comprises old key column values, new key column values, new non-key column values, an old monotonic number, a new monotonic number, an old copy identification, and a new copy identification of the captured change.

10 19. The method of claim 1, wherein the applying (e) comprises:

(e1) detecting a conflict for the communicated change with at least one row of a

target table copy;

(e2) determining a priority for the communicated change if the conflict is detected;

and

15 (e3) changing the target table copy in accordance with the detected conflict and

the determined priority for the communicated change.

20. The method of claim 19, wherein the detecting (e1) comprises:

(e1i) if the communicated change is a row delete, the conflict is detected if:

20 no row of the target table copy with key column values matching key column values

of the communicated change is identified, or

for a row of the target table copy with key column values matching the key column

values of the communicated change:

a monotonic number of the row of the target table copy does not match the
monotonic number of the communicated change, or
a copy identification of the row of the target table copy does not match a copy
identification of the communicated change.

5

21. The method of claim 19, wherein the detecting (e1) further comprises:
(e1i) if the communicated change is a row insert, the conflict is detected if:
a row of the target table copy with key column values matching key column values of
the communicated change is identified.

10

22. The method of claim 19, wherein the detecting (e1) further comprises:
(e1i) if the communicated change is a non-key update, the conflict is detected if:
no rows of the target table copy with key column values matching key column values
of the communicated change is identified, or
for a row of the target table copy with key column values matching the key column
values of the communicated change:
a monotonic number of the row of the target table copy does not match the
old monotonic number of the communicated change, or
a copy identification of the row of the target table copy does not match an old
copy identification of the communicated charge.

20

23. The method of claim 19, wherein the detecting (e1) further comprises;

(e1i) if the communicated change is a key update, the conflict is detected if:

- no row of the target table copy with key column values matching old key column values of the communicated change is identified, or
- for a row of the target table copy with key column values matching the old key column values of the communicated change:
- 5 a monotonic number of the row of the target table copy does not match an old monotonic number of the communication change, or
- a copy identification of the row of the target table copy does not match an old copy identification for the communicated change, or
- a row of the target table copy with key column values matching new key column values of the communicated change is identified.
- 10
24. The method of claim 19, wherein the determining (e2) comprises:
- (e2i) assigning priority to the communicated change, if no conflict is detected.
- 15
25. The method of claim 19, wherein the determining (e2) comprises:
- (e2i) assigning priority to the communicated change if the communication change is a row insert, and:
- a monotonic number of the communicated change is greater than a monotonic number of the conflicting row in the target table copy, or
- 20 the monotonic number of the communicated change is equal to the monotonic number of the conflicting row, and a copy identification of the communication change is greater than a copy identification of the conflicting row.

26. The method of claim 19, wherein the determining (e2) comprises:

(e2i) assigning priority to the communicated change if the communicated change is

a row delete and if:

no row in the target table copy matches key column values of the communicated

5 change, or

a monotonic number of the communicated change is greater than a monotonic

number of the conflicting row in the target table copy, or

the monotonic number for the communicated change is the same as the monotonic

number for the conflicting row and a copy identification of the communication change is

10 greater than a copy identification of the conflicting row.

27. The method of claim 19, wherein the determining (e2) comprises:

(e2i) assigning priority to the communicated change if the communicated change is

a non-key update and if:

15 no row in the target table copy matches key column values of the communicated

change, or

a monotonic number of the communicated change is greater than a monotonic

number of the conflicting row in the target table copy, or

the monotonic number of the communicated change is the same as the monotonic

20 number of the conflicting row and a copy identification of the communicated change is

greater than a copy identification of the conflicting row.

28. The method of claim 19, wherein the determining (e2) comprises:

(e2i) assigning priority to the communicated change if the communicated change is a key update and if:

(e2iA) no row in the target table copy matching old key column values of the communicated change is identified, and

5 (e2iA(I)) no row in the target table copy matching new key column values of the communicated change is identified, or

(e2iA(II)) a new monotonic number of the communicated change is greater than a monotonic number for the conflicting row in the target table copy with key column values matching new key column values of the communicated change, or

10 (e2iA(III)) the new monotonic number of the communicated change matches the monotonic number and a copy identification of the conflicting row in the target table copy with key column values matching new key column values of the communicated change, respectively; or

(e2iB) an old monotonic number of the communicated change is greater than the monotonic number of the conflicting row in the target table copy with key columns matching old key column values of the communicated change, or

15 the monotonic number of the communicated change matches the monotonic number of the conflicting row and an old copy identification of the communicated change is greater than the copy identification of the conflicting row with key column values matching old key column values of the communicated change, and

20 (e2iB(I)) no row in the target table copy matching new key column values in the communicated change is identified, or

(e2iB(II)) the new monotonic number of the communicated change is

greater than the monotonic number of the conflicting row with key column values matching the new key column values of the communicated change, or

(e2iB(III)) the monotonic number of the communicated change matches the monotonic number of the conflicting row and the new copy identification of the communicated change is greater than the copy identification of the conflicting row with key column values matching the new key column values of the communicated change.

29. The method of claim 19, wherein the changing (e3) comprises:

- (e3i) controlling propagation of the change applied to the target table copy;
 - (e3ii) insuring convergence of the plurality of table copies to the same state; and
 - (e3iii) installing the communicated changes into the target table copy.

30. The method of claim 29, wherein the controlling (e3i) comprises:

- (e3iA) if the communicated change is a row insert, setting a copy identification of the applied change to the target table copy to a copy identification received with the communicated change.

31. The method of claim 29, wherein the controlling (e3i) comprises:

- (e3iA) if the communicated change is a row delete, updating a row of the target table copy with key column values matching key column values of the communicated change by setting the delete label to indicate not to propagate the applied change.

32. The method of claim 29, wherein the controlling (e3i) comprises:

(e3iA) if the communicated change is a non-key or key update, setting a copy identification of the applied change to the target table copy to the copy identification received with the communicated change.

5 33. The method of claim 29, wherein the insuring (e3ii) comprises:

(e3iiA) insuring propagation of an implicit delete change in the target table copy by setting a conflict label of the propagation controls of the conflicting row when the communicated change is a row insert, row delete, or non-key update assigned priority, and a copy identification of the conflicting row is a copy identification assigned to the target table 10 copy.

10

15 34. The method of claim 29, wherein the insuring (e3ii) comprises:

(e3iiA) insuring propagation of a delete change in the target table by setting a conflict label of the propagation controls of the conflicting row with key column values matching old key column values of the communicated change, when the communicated change is a key update assigned priority, and a copy identification of the conflicting row matches a copy 20 identification assigned to the target table copy; and

15

(e3iiB) insuring propagation of a delete change in the target table copy by setting a conflict label of the propagation controls of the conflicting row with key column values matching new key column values of the communicated change, when the communicated 25 change is a key update assigned priority, and a copy identification of the conflicting row is a copy identification assigned to the target table copy.

35. The method of claim 29, wherein the insuring (e3ii) comprises:
(e3iiA) recording communicated old key column values, an old monotonic number,
and an old copy identification of a conflicting change in the target table copy in a delete
tombstone, when the communicated change is a conflicting delete or a conflicting update
with conflicting communicated old monotonic number or copy identification.
5

36. The method of claim 29, wherein the insuring (e3ii) comprises:
(e3iiA) checking for matching delete and suppressing application of the
communication change if a delete tombstone matching the new key columns, new
10 monotonic number, and new copy identification of the communicated change is found, when
the communicated change is an insert change or an update change with a conflicting insert
assigned priority.

37. The method of claim 1, further comprising:
15 (f) reporting each conflicting change.

38. The method of claim 37, wherein the reporting (f) comprises:
(f1) reporting conflicting changes of a row delete, a row insert, or a non-key
update change only when priority is not assigned to the communicated change and a copy
20 identification for a conflicting row in a target table copy is the copy identification assigned
to the target table copy.

39. The method of claim 37, wherein the reporting (f) comprises:

- (f1) reporting conflicting changes of a key update change only when priority is not assigned to the communicated change, and
- 5 a copy identification for a conflicting row in a target table copy with key column values matching old key column values for the communicated change is the copy identification assigned to the target table copy, or
- 10 a copy identification for a conflicting row with key column values matching new key column values of the communicated change is the copy identification assigned to the target table copy.
- 15 40. A database system, comprising:
- a plurality of table copies, wherein each row of the plurality of table copies are labeled with a monotonic number, a copy identification, and propagation controls;
- a mechanism for asynchronously capturing at least one labeled change to any row of any of the plurality of table copies from a database recovery log;
- 20 a plurality of message queues for communicating the captured changes to others of the plurality of table copies; and
- a mechanism for applying the communicated changes to the others of the plurality of table copies, wherein the plurality of table copies converge to a same state.
- 25 41. The system of claim 40, wherein the monotonic number comprises a non-decreasing time value, wherein the time values for each of the plurality of table copies are in a common time base.

42. The system of claim 41, wherein the monotonic number is automatically applied to a target table copy in response to a user induced change.

5 43. The system of claim 40, wherein the copy identification is assigned to each of the plurality of table copies, wherein the copy identification uniquely identified each of the plurality of table copies, wherein the copy identifications for each of the plurality of table copies have an ordering property.

10 44. The system of claim 40, wherein the copy identification is automatically applied to a changed table copy in response to a user induced change.

15 45. The system of claim 40, wherein the propagation controls comprise a delete label and a conflict label.

20 46. The system of claim 45, wherein the delete label indicates that a row delete in any of the plurality of table copies is not to be communicated to the others of the plurality of table copies.

25 47. The system of claim 45, wherein the conflict label indicates that an implicit row delete in any of the plurality of table copies is to be communicated to the others of the plurality of table copies.

30 48. A method for providing convergence of data copies in asynchronous data

replication in a database system, the database system including a plurality of nodes with a plurality of table copies, comprising:

- (a) labeling rows of the plurality of table copies with a monotonic number, a copy identification, and propagation controls;
- 5 (b) asynchronously capturing at least one labeled change to any row of any of the plurality of table copies from a database recovery log;
- (c) determining that the captured change is to be communicated to others of the plurality of table copies;
- 10 (d) communicating the captured change to the others of the plurality of table copies; and
- (e) applying the communicated change to the others of the plurality of table copies, wherein the plurality of table copies converge to a same state, wherein the applying comprises:
 - (e1) detecting a conflict for the communicated change with at least one row of a target table copy;
 - 15 (e2) determining a priority for the communicated change if the conflict is detected; and
 - (e3) changing the target table copy in accordance with the detected conflict and the determined priority for the communicated change.

20 49. A method for providing convergence of data copies in asynchronous data replication in a database system, the database system including a plurality of nodes with a plurality of table copies, comprising:

- (a) labeling rows of the plurality of table copies with a monotonic number, a copy identification, and propagation controls;
- (b) asynchronously capturing at least one labeled change to any row of any of the plurality of table copies from a database recovery log;
- 5 (c) determining that the captured change is to be communicated to others of the plurality of table copies;
- (d) communicating the captured change to the others of the plurality of table copies; and
- 10 (e) applying the communicated change to the others of the plurality of table copies, wherein the applying comprises:
- (e1) detecting a conflict for the communicated change with at least one row of a target table copy of the plurality of table copies;
- (e2) determining a priority for the communicated change if the conflict is detected; and
- 15 (e3) changing the target table copy in accordance with the detected conflict and the determined priority for the communicated change, wherein the changing comprises:
- (e3i) controlling propagation of the change applied to the target table copy;
- (e3ii) insuring convergence of the plurality of table copies to the same state; and
- 20 (e3iii) installing the communicated changes into the target table copy.

50. A computer readable medium with program instructions for providing

convergence of data copies in asynchronous data replication in a database system, the database system including a plurality of nodes with a plurality of table copies, comprising the instructions for:

- (a) labeling rows of the plurality of table copies with a monotonic number, a copy identification, and propagation controls;
- (b) asynchronously capturing at least one labeled change to any row of any of the plurality of table copies from a database recovery log;
- (c) determining that the captured change is to be communicated to others of the plurality of table copies;
- (d) communicating the captured change to the others of the plurality of table copies; and
- (e) applying the communicated change to the others of the plurality of table copies, wherein the plurality of table copies converge to a same state.

51. The medium of claim 50, wherein the monotonic number comprises a non-decreasing time value, wherein the time values for each of the plurality of table copies are in a common time base.

52. The medium of claim 51, wherein the monotonic number is automatically applied to a target table copy in response to a user induced change.

53. The medium of claim 50, wherein the copy identification is assigned to each of the plurality of table copies, wherein the copy identification uniquely identifies each of

the plurality of table copies, wherein the copy identifications for each of the plurality of table copies have an ordering property.

5 54. The medium of claim 53, wherein the copy identification is automatically applied to a changed table copy in response to a user induced change.

10 55. The medium of claim 50, wherein the propagation controls comprise a delete label and a conflict label.

15 56. The medium of claim 55, wherein the delete label indicates that a row delete in any of the plurality of table copies is not to be communicated to the others of the plurality of table copies.

15 57. The medium of claim 55, wherein the conflict label indicates that an implicit row delete in any of the plurality of table copies is to be communicated to the others of the plurality of table copies.

20 58. The medium of claim 50, wherein the capturing instruction (b) comprises instructions for:

- (b1) accessing at least one entry of the database recovery log pertaining to the labeled change;
- (b2) determining a type of the labeled change;
- (b3) extracting old column values and new column values of the labeled change

from the database recovery log entry; and

(b4) extracting changed key column values and unchanged key column values

from the database recovery log entry.

5 59. The medium of claim 58, wherein the type of the labeled change comprises a row insert, a row delete, a non-key update, or a key update.

60. The medium of claim 50, wherein the determining instruction (c) comprises instructions for:

10 (c1) examining a labeled change type, the copy identification, and the propagation control values for the captured change, wherein the propagation control values comprises a delete label and a conflict label.

15 61. The medium of claim 60, wherein the determining instruction (c) further comprises instructions for:

(c2) determining that the captured change is to be communicated to the others of the plurality of table copies if the labeled change type is a row insert and if a copy identification for the captured change is a copy identification assigned to a current table copy.

20 62. The medium of claim 60, wherein the determining instruction (c) further comprises instructions for:

(c2) determining that the captured change is to be communicated to the others of

the plurality of table copies if the labeled change type is a row delete and if the delete label of the captured change indicates that the captured change is to be communicated.

63. The medium of claim 60, wherein the determining instruction (c) further
5 comprises instructions for:

(c2) determining that the captured change is to be communicated to the others of the plurality of table copies if the labeled change type is a non-key update or a key update and if neither the delete label nor the conflict label of the captured change indicates that the captured change is not to be communicated, and if a new copy identification of the captured
10 change is a copy identification assigned to a current table copy.

64. The medium of claim 50, wherein the communicating instruction (d)
comprises instructions for:

(d1) sending information for the captured change to the others of the plurality of
15 table copies, wherein if the captured change is a row insert, the information comprises key column values, non-key column values, a monotonic number, and a copy identification of the captured change.

65. The medium of claim 50, wherein the communicating instruction (d)
20 comprises instructions for:

(d1) sending information for the captured change to the others of the plurality of table copies, wherein if the captured change is a row delete, the information comprises key column values, a monotonic number, and a copy identification of the captured change.

66. The medium of claim 50, wherein the communicating instruction (d) comprises instructions for:

(d1) sending information for the captured change to the others of the plurality of table copies, wherein if the captured change is a non-key update, the information comprises key column values, new non-key column values, an old monotonic number, a new monotonic number, an old copy identification, and a new copy identification of the captured change.

5
67. The medium of claim 50, wherein the communicating instruction (d) comprises instructions for:

10
(d1) sending information for the captured change to the others of the plurality of table copies, wherein if the captured change is a key update, the information comprises old key column values, new key column values, new non-key column values, an old monotonic number, a new monotonic number, an old copy identification, and a new copy identification of the captured change.

15
68. The medium of claim 50, wherein the applying instruction (e) comprises instructions for:

20
(e1) detecting a conflict for the communicated change with at least one row of a target table copy;

(e2) determining a priority for the communicated change if the conflict is detected; and

(e3) changing the target table copy in accordance with the detected conflict and

the determined priority for the communicated change.

69. The medium of claim 68, wherein the detecting instruction (e1) comprises instructions for:

5 (e1i) if the communicated change is a row delete, the conflict is detected if:

no row of the target table copy with key column values matching key column values of the communicated change is identified, or

for a row of the target table copy with key column values matching the key column values of the communicated change:

10 a monotonic number of the row of the target table copy does not match the monotonic number of the communicated change, or

a copy identification of the row of the target table copy does not match a copy identification of the communicated change.

15 70. The medium of claim 68, wherein the detecting instruction (e1) further comprises instructions for:

(e1i) if the communicated change is a row insert, the conflict is detected if:

a row of the target table copy with key column values matching key column values of the communicated change is identified.

20 71. The medium of claim 68, wherein the detecting instruction (e1) further comprises instructions for:

(e1i) if the communicated change is a non-key update, the conflict is detected if:

no rows of the target table copy with key column values matching key column values of the communicated change is identified, or

for a row of the target table copy with key column values matching the key column values of the communicated change:

5 a monotonic number of the row of the target table copy does not match the old monotonic number of the communicated change, or
a copy identification of the row of the target table copy does not match an old copy identification of the communicated change.

10 72. The medium of claim 68, wherein the detecting instruction (e1) further comprises instructions for:

(e1i) if the communicated change is a key update, the conflict is detected if:
no row of the target table copy with key column values matching old key column values of the communicated change is identified, or
for a row of the target table copy with key column values matching the old key column values of the communicated change:

a monotonic number of the row of the target table copy does not match an old monotonic number of the communication change, or
a copy identification of the row of the target table copy does not match an old copy identification for the communicated change, or
a row of the target table copy with key column values matching new key column values of the communicated change is identified.

73. The medium of claim 68, wherein the determining instruction (e2) comprises instructions for:

(e2i) assigning priority to the communicated change, if no conflict is detected.

5 74. The medium of claim 68, wherein the determining instruction (e2) comprises instructions for:

(e2i) assigning priority to the communicated change if the communication change is a row insert, and:

a monotonic number of the communicated change is greater than a monotonic
10 number of the conflicting row in the target table copy, or
the monotonic number of the communicated change is equal to the monotonic
number of the conflicting row, and a copy identification of the communication change is
greater than a copy identification of the conflicting row.

15 75. The medium of claim 68, wherein the determining instruction (e2) comprises instructions for:

(e2i) assigning priority to the communicated change if the communicated change is a row delete and if:

20 no row in the target table copy matches key column values of the communicated
change, or
a monotonic number of the communicated change is greater than a monotonic
number of the conflicting row in the target table copy, or
the monotonic number for the communicated change is the same as the monotonic

number for the conflicting row and a copy identification of the communicated change is greater than a copy identification of the conflicting row.

5 76. The medium of claim 68, wherein the determining instruction (e2) comprises

instructions for:

(e2i) assigning priority to the communicated change if the communicated change is a non-key update and if:

no row in the target table copy matches key column values of the communicated change, or

10 a monotonic number of the communicated change is greater than a monotonic number of the conflicting row in the target table copy, or
the monotonic number of the communicated change is the same as the monotonic number of the conflicting row and a copy identification of the communicated change is greater than a copy identification of the conflicting row.

15

77. The medium of claim 68, wherein the determining instruction (e2) comprises instructions for:

(e2i) assigning priority to the communicated change if the communicated change is a key update and if:

20 (e2iA) no row in the target table copy matching old key column values of the communicated change is identified, and

(e2iA(I)) no row in the target table copy matching new key column values of the communicated change is identified, or

(e2iA(II)) a new monotonic number of the communicated change is greater than a monotonic number for the conflicting row in the target table copy with key column values matching new key column values of the communicated change, or

(e2iA(III)) the new monotonic number and a copy identification of the communicated change matches the monotonic number and a copy identification of the conflicting row in the target table copy with key column values matching new key column values of the communicated change, respectively; or

(e2iB) an old monotonic number of the communicated change is greater than the monotonic number of the conflicting row in the target table copy with key columns matching old key column values of the communicated change, or

the monotonic number of the communicated change matches the monotonic number of the conflicting row and an old copy identification of the communicated change is greater than the copy identification of the conflicting row with key column values matching old key column values of the communicated change, and

(e2iB(I)) no row in the target table copy matching new key column values in the communicated change is identified, or

(e2iB(II)) the new monotonic number of the communicated change is greater than the monotonic number of the conflicting row with key column values matching the new key column values of the communicated change, or

(e2iB(III)) the monotonic number of the communicated change matches the monotonic number of the conflicting row and the new copy identification of the communicated change is greater than the copy identification of the conflicting row with key column values matching the new key column values of the communicated change.

78. The medium of claim 68, wherein the changing instruction (e3) comprises instructions for:

- (e3i) controlling propagation of the change applied to the target table copy;
- (e3ii) insuring convergence of the plurality of table copies to the same state; and
- (e3iii) installing the communicated changes into the target table copy.

5
79. The medium of claim 78, wherein the controlling instruction (e3i) comprises instructions for:

10 (e3iA) if the communicated change is a row insert, setting a copy identification of the applied change to the target table copy to a copy identification received with the communicated change.

80. The medium of claim 78, wherein the controlling instruction (e3i) comprises instructions for:

15 (e3iA) if the communicated change is a row delete, updating a row of the target table copy with key column values matching key column values of the communicated change by setting the delete label to indicate not to propagate the applied change.

20
81. The medium of claim 78, wherein the controlling instruction (e3i) comprises instructions for:

(e3iA) if the communicated change is a non-key or key update, setting a copy identification of the applied change to the target table copy to the copy identification received with the communicated change.

82. The medium of claim 78, wherein the insuring instruction (e3ii) comprises instructions for:

(e3iiA) insuring propagation of an implicit delete change in the target table copy by setting a conflict label of the propagation controls of the conflicting row when the communicated change is a row insert, a row delete, or a non-key update assigned priority and a copy identification of the conflicting row is a copy identification assigned to the target table copy

5
83. The medium of claim 78, wherein the insuring instruction (e3ii) comprises instructions for:

(e3iiA) insuring propagation of a delete change in the target table by setting a conflict label of the propagation controls of the conflicting row with key column values matching old key column values of the communicated change, when the communicated change is a key update assigned priority and a copy identification of the conflicting row matches a copy identification assigned to the target table copy.

10
15
84. The medium of claim 78, wherein the insuring instruction (e3ii) comprises instructions for:

(e3iiA) recording communicated old key column values, an old monotonic number, and an old copy identification of a conflicting change in the target table copy in a delete tombstone, when the communicated change is a conflicting delete or a conflicting update with conflicting communicated old monotonic number or copy identification.

85. The medium of claim 78, wherein the insuring instruction (e3ii) comprises instructions for:

(e3iiA) checking for matching delete and suppressing application of the communication change if a delete tombstone matching new key column values, new monotonic number, and new copy identification of the communicated change is found, when the communicated change is an insert change or an update change with a conflicting insert assigned priority.

5
86. The medium of claim 50, further comprising instructions for:

10 (f) reporting each conflicting change.

15
87. The medium of claim 86, wherein the reporting instruction (f) comprises instructions for:

(f1) reporting conflicting changes of a row delete, a row insert, or a non-key update change only when priority is not assigned to the communicated change and a copy identification for a conflicting row in a target table copy is the copy identification assigned to the target table copy.

20
88. The medium of claim 86, wherein the reporting instruction (f) comprises instructions for:

(f1) reporting conflicting changes of a key update change only when priority is not assigned to the communicated change, and a copy identification for a conflicting row in a target table copy with key column

values matching old key column values for the communicated change is the copy identification assigned to the target table copy, or
5 a copy identification for a conflicting row with key column values matching new key column values of the communicated change is the copy identification assigned to the target table copy.

89. A computer readable medium with program instructions for providing convergence of data copies in asynchronous data replication in a database system, the database system including a plurality of nodes with a plurality of table copies, comprising
10 the instructions for:

- (a) labeling rows of the plurality of table copies with a monotonic number, a copy identification, and propagation controls;
- (b) asynchronously capturing at least one labeled change to any row of any of the plurality of table copies from a database recovery log;
- (c) determining that the captured change is to be communicated to others of the plurality of table copies;
- (d) communicating the captured change to the others of the plurality of table copies; and
- (e) applying the communicated change to the others of the plurality of table copies, wherein the plurality of table copies converge to a same state, wherein the applying comprises:
 - (e1) detecting a conflict for the communicated change with at least one row of a target table copy;

(e2) determining a priority for the communicated change if the conflict is detected; and

(e3) changing the target table copy in accordance with the detected conflict and the determined priority for the communicated change.

5

90. A computer readable medium with program instructions for providing convergence of data copies in asynchronous data replication in a database system, the database system including a plurality of nodes with a plurality of table copies, comprising the instructions for:

- 10 (a) labeling rows of the plurality of table copies with a monotonic number, a copy identification, and propagation controls;
- (b) asynchronously capturing at least one labeled change to any row of any of the plurality of table copies from a database recovery log;
- (c) determining that the captured change is to be communicated to others of the plurality of table copies;
- 15 (d) communicating the captured change to the others of the plurality of table copies; and
- (e) applying the communicated change to the others of the plurality of table copies, wherein the applying comprises:

20 (e1) detecting a conflict for the communicated change with at least one row of a target table copy of the plurality of table copies;

(e2) determining a priority for the communicated change if the conflict is detected; and

- (e3) changing the target table copy in accordance with the detected conflict and the determined priority for the communicated change, wherein the changing comprises:
- (e3i) controlling propagation of the change applied to the target table copy;
- 5 (e3ii) insuring convergence of the plurality of table copies to the same state; and
- (e3iii) installing the communicated changes into the target table copy.